



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Scutari, M;Mackay, I;Balding, D

Title:

Improving the efficiency of genomic selection

Date:

2013-08-01

Citation:

Scutari, M., Mackay, I. & Balding, D. (2013). Improving the efficiency of genomic selection. *Statistical Applications in Genetics and Molecular Biology*, 12 (4), pp.517-527. <https://doi.org/10.1515/sagmb-2013-0002>.

Persistent Link:

<https://hdl.handle.net/11343/52576>

Improving the Efficiency of Genomic Selection

Marco Scutari¹, Ian Mackay², and David Balding¹

¹Genetics Institute, University College London (UCL), United Kingdom

²National Institute of Agricultural Botany (NIAB), Cambridge, United kingdom

June 4, 2013

Abstract

We investigate two approaches to increase the efficiency of phenotypic prediction from genome-wide markers, which is a key step for genomic selection (GS) in plant and animal breeding. The first approach is feature selection based on Markov blankets, which provide a theoretically-sound framework for identifying non-informative markers. Fitting GS models using only the informative markers results in simpler models, which may allow cost savings from reduced genotyping. We show that this is accompanied by no loss, and possibly a small gain, in predictive power for four GS models: partial least squares (PLS), ridge regression, LASSO and elastic net. The second approach is the choice of kinship coefficients for genomic best linear unbiased prediction (GBLUP). We compare kinships based on different combinations of centring and scaling of marker genotypes, and a newly proposed kinship measure that adjusts for linkage disequilibrium (LD).

We illustrate the use of both approaches and examine their performances using three real-world data sets with continuous phenotypic traits from plant and animal genetics. We find that elastic net with feature selection and GBLUP using LD-adjusted kinships performed similarly well, and were the best-performing methods in our study.

Keywords: genome-wide prediction, genomic selection, feature selection, Markov blanket, linkage disequilibrium, kinship.

1 Introduction

The ever-increasing amount of genetic information available in plant and animal breeding is reflected in the development of sophisticated models for the prediction of quantitative traits from genome-wide markers (Heffner et al., 2009, Hayes et al., 2009), also known as genomic selection (GS). The markers are typically dense single-nucleotide polymorphisms (SNPs). Approaches to this problem have moved from models with simple specifications, such as ridge regression (Hoerl and Kennard, 1970) and the LASSO (Tibshirani, 1996), to models based on highly-structured hierarchical distributions or semiparametric approaches. Some examples are the Bayesian alphabet models (Gianola et al., 2009, de los Campos et al., 2012), Bayesian models with complex priors as in Guan and Stephens (2011), models based on reproducing kernel Hilbert spaces (RKHS) such as Bravo et al. (2009), or the Bayesian LASSO (Park and Casella, 2008, Vazquez et al., 2012).

This complexity is motivated by the need to correctly model the genetic architecture of the trait under investigation while producing models that are easy to estimate even for large SNP profiles. We focus on two key aspects of these models: the inclusion of a preliminary step that removes SNPs that appear to be redundant, and the choice of kinship matrices to model the relatedness of the genotyped individuals.

The former is equivalent to *feature selection* (Koller and Sahami, 1996), and can also be achieved by shrinking SNP effects towards zero, either through the use of constraints (Zou and Hastie, 2005) or through appropriate prior distributions in a Bayesian setting (Meuwissen et al., 2001). We examine the effectiveness in GS of Markov blankets (Pearl, 1988), which have been extensively studied in graphical modelling. They provide a principled solution to feature selection problems, and can be implemented as a data pre-processing step prior to fitting the GS model. We implement Markov blanket feature selection within four GS models applied to three real-world data sets covering barley, rice and mouse genetics.

Kinship matrices were traditionally derived from pedigrees using a single definition, but with kinships now being calculated from SNP data many different definitions are available (Astle and Balding, 2009). We investigate four kinship matrices within genetic best linear unbiased prediction (GBLUP). These include a novel matrix introduced by Speed et al. (2012a) which adjusts for the bias introduced by differences in local linkage disequilibrium (LD), and has been shown to increase the precision of heritability estimates.

2 Background

2.1 Markov Blankets and Feature Selection

The Markov Blanket of a variable of interest T , denoted as $B(T)$, is the minimal set of variables conditioned on which all other variables in the model are probabilistically independent of the target T (Pearl, 1988). The Markov blanket of a phenotype \mathbf{y} in a GS model is the minimal set $B(\mathbf{y}) \subset \mathbf{X}$ such that

$$P(\mathbf{y} | \mathbf{X}) = P(\mathbf{y} | B(\mathbf{y})), \quad (1)$$

that is, the subset of SNPs $B(\mathbf{y})$ that makes all other SNPs redundant as far as the trait \mathbf{y} is concerned. Given this property, knowledge of only the SNPs in $B(\mathbf{y})$ is enough to determine the probability distribution of \mathbf{y} . Other SNPs become superfluous, either because they are not associated with the trait or because their effect is mediated by the SNPs in $B(\mathbf{y})$. If $B(\mathbf{y})$ were known, any GS model could be fitted using $B(\mathbf{y})$ instead of the full SNP profile \mathbf{X} with no loss of information, but in practice the need to estimate $B(\mathbf{y})$ means that some information loss is possible. This two-stage approach contrasts with models such as BayesB (Meuwissen et al., 2001) and the LASSO (Tibshirani, 1996), which select significant SNP effects concurrently with model fitting and in a model-specific way.

Markov blankets can be efficiently estimated from data through the use of conditional independence tests, such as parametric and non-parametric tests for partial correlation (Legendre, 2000, Hotelling, 1953) or mutual information (Scutari and Brogini, 2012). Tests in common use do not require any tuning parameter except for the type I error threshold α . The estimated $B(\mathbf{y})$ will satisfy (1) only approximately because of type I and type II errors. The former arise from the noisiness inherent to the data and limited sample sizes, while the latter are typical of weak dependencies which will often be omitted from the Markov blanket.

Several computationally-efficient heuristic algorithms for Markov blanket estimation are available in literature, including Grow-Shrink (GS; Margaritis, 2003), Incremental Association (IAMB; Tsamardinos et al., 2003) and Hiton-MB (Aliferis et al., 2010). For instance, IAMB can be used to estimate the Markov blanket of a trait \mathbf{y} as follows:

1. Set $B(\mathbf{y}) = \{\emptyset\}$.
2. **Forward Phase:** until no change is made,
 - (a) test each SNP X_i for independence from \mathbf{y} conditional on the current Markov blanket $B(\mathbf{y})$;
 - (b) admit into $B(\mathbf{y})$ the SNP whose test returned the lowest p-value if that p-value is smaller than α .

3. **Backward Phase:** for each $X_i \in B(\mathbf{y})$, remove X_i from $B(\mathbf{y})$ if \mathbf{y} is independent of X_i conditional on $B(\mathbf{y}) \setminus X_i$.

As a result, conditional independence tests are performed in order of increasing complexity, thus ensuring that in practice only a small number of SNPs is used for each test. Compared to single-SNP analyses, such as those described in Macciotta et al. (2009) and Schulz-Streeck et al. (2011), feature selection with Markov blankets is computationally more expensive because of the use of conditional ($\mathbf{y} \perp\!\!\!\perp X_i | B(\mathbf{y}) \setminus X_i$) instead of marginal ($\mathbf{y} \perp\!\!\!\perp X_i$) independence tests. However, as shown in Section 4, Figure 5, conditional tests are more effective at discarding SNPs that carry essentially the same information about the trait and select subsets with more predictive power for the same size.

2.2 Kinship Estimation

In the past, pedigree information was used to specify kinships, but such information is often missing or inadequate. SNP-based methods for measuring kinships have become increasingly common and have the advantage of measuring the realised amount of genome sharing, as opposed to the expected value provided by pedigree-based methods (Astle and Balding, 2009, Forni et al., 2011)

The SNP-based kinship of two individuals is usually based on the average over SNPs of the product of their genotypes, coded as 0, 1 and 2 according to the count of one of the two alleles. By design, it can only capture the additive components of kinship, and it has very low power in identifying non-additive ones. In the following, we denote this genotype matrix with \mathbf{X} , with rows corresponding to individuals and columns to SNPs, and with X_i its i th column.

In human genetics, kinship is commonly measured as the proportion of shared alleles at each locus (Morris and Cardon, 2007). This approach is also known as identical-by-state (IBS) kinship, and will be denoted by \mathbf{K}_0 . Unlike other kinship matrices below, \mathbf{K}_0 is always non-negative. However, it cannot be expressed in the form $\mathbf{X}\mathbf{X}^T$, which leads to parameters directly interpretable as SNP effect sizes (see Section 3 for details).

Another choice, common in plant and animal genetics, is to centre the genotypes (Habier et al., 2007, VanRaden, 2008) and estimate the kinship matrix as

$$\mathbf{K}_1 = \frac{1}{m} \sum_{i=1}^m (X_i - 2p_i)(X_i - 2p_i)^T \quad (2)$$

where m is the number of markers and p_i is a vector with every entry equal to the population allele fraction, usually estimated as the mean of $X_i/2$. Centring improves interpretability, since kinship values can be interpreted as an excess or deficiency of allele sharing compared with random allocation of alleles,

and so zero can be interpreted as “unrelated”. However, the requirement to estimate the p_i , usually from the same data set, can cause problems in some settings (Astle and Balding, 2009).

One criticism of both the above choices is that the sharing of a rare allele between two individuals counts the same as the sharing of a common allele. One natural approach to giving more weight to the sharing of a rare allele is to standardise over SNPs, thus obtaining

$$\mathbf{K}_2 = \frac{1}{m} \sum_{i=1}^m \bar{X}_i \bar{X}_i^T \quad \text{where} \quad \bar{X}_i = \frac{X_i - 2p_i}{\sqrt{2p_i(1-p_i)}}. \quad (3)$$

The (i, j) entry of \mathbf{K}_2 can be interpreted as an average over SNPs of the correlation coefficient estimated from a single pair of individuals, i and j (Astle and Balding, 2009).

A modification of \mathbf{K}_2 has been recently proposed by Speed et al. (2012a), based on evidence that the effects of SNPs are sensitive to uneven LD across the genome. In particular, SNP effects are over-estimated in high-LD regions and under-estimated in low-LD regions due the uneven tagging of causal variants. The contributions of causal variants are picked up by a larger number of SNPs in high-LD regions compared to low-LD regions, thus introducing bias in the GS models and in turn in subsequent inference such as prediction or heritability estimation. To correct for this bias, SNPs can be re-weighted:

$$\mathbf{K}_3 = \frac{\sum_{i=1}^m w_i \bar{X}_i \bar{X}_i^T}{\sum_{i=1}^m w_i}, \quad (4)$$

where the weight vector $\mathbf{w} = [w_1 \cdots w_m]$ solves

$$\min_{\mathbf{w}} \sum_{i=1}^m |\mathbf{1} - \mathbf{C}_i \mathbf{w}| \quad \text{subject to} \quad w_1, \dots, w_m > 0 \quad (5)$$

and \mathbf{C}_i is a vector of squared correlations of SNP i with neighbouring SNPs. SNP effects are set to decay exponentially with physical distance, according to a decay rate λ whose value reflects the average LD for the data set. As a result we have that

$$w_i + \sum_{j \neq i} w_j C_{ij} e^{-\lambda d_{ij}}, \quad (6)$$

where d_{ij} is the distance between SNPs i and j , is approximately constant as the weights offset differences in LD as measured by the squared correlations C_{ij} .

For computational reasons, the minimisation in (5) is performed separately on different chromosomes and, within each chromosome, on different regions chosen based on λ .

3 Materials and Methods

We explored the effects of the approaches outlined in Section 2 on the predictive power of GS models using three publicly-available real-world data sets including continuous phenotypic traits. The yield data from the AGOUEB project (Waugh et al., 2010, Cockram et al., 2010) consist of 227 UK winter barley varieties and 810 SNPs. The heterogeneous mouse population (Solberg et al., 2006, Valdar et al., 2006) from the Wellcome Trust Case Control Consortium (WTCCC) consists of 1940 SNP profiles and 12545 SNPs; among the recorded traits, we consider growth rate and weight. The rice data set from Zhao et al. (2011) consists of 413 varieties of *Oryza sativa* with 73808 SNPs; among the 34 recorded traits, we consider the number of seeds per panicle because of its low variability among the various subpopulations included in the original analysis.

All data sets have been preprocessed by removing SNPs with minor allele frequencies $< 1\%$ and those with $> 20\%$ missing data. The missing data in the remaining SNPs have been imputed using the **impute** R package (Hastie et al., 2012). Other widely used imputation methods in genetics, such as that implemented in MaCH (Li et al., 2010), were not available because of the absence of accurate SNP maps; the position of many SNPs is unknown, and only genetic distances (in cM) were available between mapped SNPs. Furthermore, we removed one SNP from each pair whose allele counts have correlation > 0.90 to increase the numerical stability of the models.

To investigate Markov blanket feature selection, we considered the following GS models:

- Ridge regression, LASSO and the elastic net penalised regressions. These are all based on

$$\mathbf{y} = \mu + \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon} \quad \text{with} \quad \hat{\boldsymbol{\beta}} = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} \{ \lambda_1 \|\boldsymbol{\beta}\|_1 + \lambda_2 \|\boldsymbol{\beta}\|_2 \}, \lambda_1, \lambda_2 \geq 0, \quad (7)$$

where \mathbf{y} is the trait of interest, \mathbf{X} are the SNP genotypes, $\boldsymbol{\beta}$ are the fixed SNP effects and $\boldsymbol{\varepsilon}$ are independent, normally-distributed errors with variance $\sigma_{\boldsymbol{\varepsilon}}^2$. We used the implementations provided by the **penalized** (Goeman, 2012) and **glmnet** (Friedman et al., 2010) R packages. When considering the elastic net we restricted both the L_1 and L_2 penalties for the genetic effects $\boldsymbol{\beta}$ to be strictly positive ($\lambda_1, \lambda_2 > 0$), to facilitate the comparison with ridge regression ($\lambda_1 = 0$) and the LASSO ($\lambda_2 = 0$).

- Partial least squares (PLS) regression as implemented in the **pls** R package (Mevik et al., 2011).

- Genetic BLUP (GBLUP) implemented in the **synbreed** R package (Wimmer et al., 2012). It uses the linear mixed model

$$\mathbf{y} = \mu + \mathbf{Z}\mathbf{g} + \varepsilon, \quad \mathbf{g} \sim N(\mathbf{0}, \mathbf{K}\sigma_g^2), \quad (8)$$

where \mathbf{g} are the random effects and \mathbf{Z} is a design matrix that can be used for example to indicate the same genotype exposed to different environments. Any positive definite matrix can be used for \mathbf{K} . Fixed effects can also be included in (8) in order to capture purely environmental effects (Heffner et al., 2009).

When \mathbf{K} can be expressed in the form $\mathbf{X}\mathbf{X}^T$, GBLUP can be shown to be equivalent to the Bayesian linear regression

$$\mathbf{y} = \sum_{i=1}^m X_i^* \beta_i + \varepsilon \quad \text{with SNP effect prior} \quad \beta_i \sim N\left(\mathbf{0}, \frac{\sigma_g^2}{m} \mathbf{I}\right), \quad (9)$$

in which \mathbf{K} determines the transformation X^* of the SNP genotypes. For instance, the original X_i are used when $\mathbf{K} = \mathbf{K}_1$; the scaled \bar{X}_i from (3) when $\mathbf{K} = \mathbf{K}_2$; and the weighted $w_i \bar{X}_i / \sum w_i$ from (4) when $\mathbf{K} = \mathbf{K}_3$. This formulation of GBLUP results in a more natural interpretation of SNP effects, and is sometimes known as random regression BLUP (RR-BLUP). An overview of its properties can be found in Piepho (2009) and Piepho et al. (2012).

Markov blanket feature selection has been performed with the IAMB algorithm as implemented in the **bnlearn** R package (Scutari, 2010), using the exact Student's t test for Pearson's correlation with a type I error threshold of $\alpha = 0.15$. Each GS model was fitted both using all the available SNPs and using only the SNPs included in the Markov blanket.

The different kinship matrices were investigated within GBLUP, as the other GS models do not include an explicit kinship term. \mathbf{K}_1 and \mathbf{K}_2 were computed using **synbreed**. For \mathbf{K}_3 , we used the freely available LDAK software (Speed et al., 2012b). The LD decay rate was set to $\lambda = 50\text{cM}$ for the AGOUEB data, $\lambda = 0.2\text{cM}$ for the mouse data and $\lambda = 100\text{cM}$ for the rice data. Such values were found, through experimentation, to ensure the LD adjustment was effective without affecting the genetic information present in the SNP profiles. \mathbf{K}_0 was computed with PLINK (Purcell et al., 2007). All configurations of GS models and kinships were fitted once using all SNPs available after preprocessing the data and once using only those in the Markov blanket.

The predictive power of the GS models was measured with Pearson's correlation coefficient ρ between the observed trait values and the predictions obtained from 10-fold cross-validation. For each model, cross-validation was run 5 times.

Markov blankets, kinship matrices and GS models were fitted separately for each fold in each cross-validation run, and the resulting correlations averaged. The correlation between observed and fitted trait values is also reported as a measure of goodness of fit.

4 Results

Table 1 reports the observed correlations (ρ , i.e. the correlation between the observed and the fitted trait values) and the predictive correlations (ρ_{CV} , i.e. the correlations obtained from cross-validation) for PLS, ridge regression, LASSO and the elastic net. The corresponding correlations arising from the subset of SNPs included in the Markov blankets are labelled ρ_{MB} and $\rho_{CV,MB}$, respectively.

First of all, we note that for $\alpha = 0.15$ Markov blankets only select a small number of SNPs, regardless of the dimension of the SNP profile. The average size of the Markov blankets obtained from cross-validation is 185 for the AGOUEB data, 543 (for growth rate) and 525 (for weight) for the mouse data, and 293 for the rice data. Of those SNPs, 136 (74%) appear in at least half of the cross-validation folds for AGOUEB, 241 (46%) for the mouse data and weight, 276 (51%) for the mouse data and growth rate, but only 15 (5%) for the rice data. This can be attributed to the very low ratio between sample size and number of SNPs in the rice data (< 0.01) compared to the mouse (0.15) and AGOUEB (0.28) data. As expected, the dimension reduction is smaller in the case of the AGOUEB data because of the limited number of available SNPs, despite the extensive LD present in cultivated UK barley (Cockram et al., 2010, Rostoks et al., 2006). On the other hand, only a small proportion of the original SNPs are retained for the mouse and rice data sets (about 4% and 0.4%, respectively). In each case, the number of SNPs included in the Markov blankets is smaller than the sample size, thus ensuring the regularity and numerical stability of the GS models.

The position of mapped SNPs within the respective genomes is shown in Figure 1 (AGOUEB), Figure 2 (rice), Figure 3 (mice, weight) and Figure 4 (mice, growth). For all but the AGOUEB data, we can see how the Markov blankets arising from cross-validation identify some regions as associated with the trait being modelled (e.g. SNPs in the range [0.1cM, 75.8cM] of chromosome 1 are included with high probability for both traits in the mice data set) while completely discarding other regions (e.g. [96.3cM, 108.3cM] in chromosome 2 and [70.6cM, 88.5cM] in chromosome 3). The positions of these regions may provide useful prior information in subsequent association studies and in targeting future genotyping efforts. In the case of the AGOUEB data, marker density is not high enough to identify regions with markedly different association levels.

Model	ρ	ρ_{MB}	Δ_1	ρ_{CV}	$\rho_{CV,MB}$	Δ_2
AGOUEB, YIELD (227 obs., 185 SNPs out of 810, 23%)						
PLS	0.812	0.805	-0.007	0.495	0.495	+0.000
Ridge	0.817	0.765	-0.051	0.501	0.489	-0.012
LASSO	0.829	0.811	-0.018	0.400	0.399	-0.001
Elastic Net	0.806	0.752	-0.054	0.500	0.489	-0.011
MICE, GROWTH RATE (1940 obs., 543 SNPs out of 12.5K, 4%)						
PLS	0.716	0.882	+0.166	0.344	0.388	+0.044
Ridge	0.841	0.889	+0.047	0.366	0.394	+0.028
LASSO	0.717	0.881	+0.164	0.390	0.394	+0.004
Elastic Net	0.751	0.893	+0.142	0.403	0.401	-0.001
MICE, WEIGHT (1940 obs., 525 SNPs out of 12.5K, 4%)						
PLS	0.927	0.823	-0.104	0.502	0.524	+0.022
Ridge	0.877	0.843	-0.034	0.526	0.542	+0.016
LASSO	0.743	0.807	+0.064	0.579	0.577	-0.001
Elastic Net	0.789	0.845	+0.056	0.580	0.580	+0.000
RICE, SEEDS PER PANICLE (413 obs., 293 SNPs out of 74K, 0.4%)						
PLS	0.853	0.923	+0.070	0.583	0.601	+0.018
Ridge	0.950	0.921	-0.029	0.601	0.612	+0.011
LASSO	0.885	0.939	+0.054	0.516	0.580	+0.064
Elastic Net	0.958	0.917	+0.040	0.602	0.612	+0.010

Table 1: Correlation coefficients for PLS, ridge regression, LASSO and the elastic net: ρ is the correlation between observed and fitted trait values; ρ_{CV} is the predictive correlation obtained from cross-validation; ρ_{MB} and $\rho_{CV,MB}$ are the corresponding quantities obtained using only the SNPs in the Markov blanket. $\Delta_1 = \rho_{MB} - \rho$ and $\Delta_2 = \rho_{CV,MB} - \rho_{CV}$. The highest value for each quantity and data set is shown in bold. The average dimension of the Markov blanket over cross-validation is reported in parentheses for each data set and trait.




Figure 1: Frequency of the SNPs included in the Markov blankets estimated from the AGOUEB data using cross-validation, plotted against the position of the SNPs in the barley genome. Green ticks indicate the positions of all mapped SNPs for this data set.




Figure 2: Frequency of the SNPs included in the Markov blankets estimated from the rice data using cross-validation, plotted against the position of the SNPs in the genome. Green ticks indicate the positions of all mapped SNPs for this data set.




Figure 3: Frequency of the SNPs included in the Markov blankets estimated from the mouse weight data using cross-validation, plotted against the position of the SNPs in the barley genome. Green ticks indicate the positions of all mapped SNPs for this data set.




Figure 4: Frequency of the SNPs included in the Markov blankets estimated from the mouse growth data using cross-validation, plotted against the position of the SNPs in the barley genome. Green ticks indicate the positions of all mapped SNPs for this data set.

Kinship matrix	AGOUEB		MICE, GROWTH		MICE, WEIGHT		RICE	
	ρ	ρ_{CV}	ρ	ρ_{CV}	ρ	ρ_{CV}	ρ	ρ_{CV}
\mathbf{K}_0	0.848	0.511	0.838	0.376	0.931	0.536	0.933	0.596
\mathbf{K}_1	0.847	0.512	0.656	0.366	0.882	0.507	0.933	0.590
\mathbf{K}_2	0.848	0.513	0.688	0.388	0.883	0.508	0.933	0.598
\mathbf{K}_3	0.832	0.521	0.695	0.400	0.881	0.554	0.918	0.594

Table 2: Correlation coefficients obtained in GBLUP using the four kinship matrices defined in Section 2.2. The highest value for each quantity and data set is shown in bold. ρ and ρ_{CV} are defined as in Table 1.

We observe no loss in the predictive power of the GS models after the Markov blanket feature selection. In fact, the increased numerical stability resulting from the reduced number of SNPs slightly improved the predictive power of the GS models. The average of ρ_{CV} over the four analyses was 0.481, 0.498, 0.471 and 0.521 for PLS, ridge, LASSO and elastic net respectively, while the corresponding averages for $\rho_{CV,MB}$ were 0.502, 0.509, 0.487 and 0.520, all with an approximate standard deviation of 0.0057 computed as in Hooper (1958).

If we choose $\alpha < 0.15$, we obtain Markov blankets that are too small to capture polygenic effects (results not shown). A possible explanation for this behaviour may be that large values of α allow Markov blankets to initially include SNPs that are weakly associated with the trait, to the point that they would be individually discarded. In addition, among them there may be sets of SNPs that are jointly significant due to epistasis, and such sets are retained in the Markov blanket.

Furthermore, Markov blankets outperform other subsamples of the same size. To show this, we generated for each data set 100 random subsets of SNPs of the same size as the corresponding Markov blanket. In addition, we also generated subsamples including the most significant SNPs from a single-SNP analysis under cross-validation. The same t test as in Markov blanket estimation was used to assess significance. Subsequently, we used them to fit the GS models and to compute the predictive correlations corresponding $\rho_{CV,MB}$. As we can see from Figure 5, the Markov blanket always results in higher values of $\rho_{CV,MB}$.

The elastic net consistently outperforms the other GS models both with and without the use of Markov blankets, except for the AGOUEB data set (in which ρ_{CV} is essentially the same for ridge regression and the elastic net).

Overall, from Table 2 we see that the predictive performance of GBLUP improves as the kinship matrices progress from \mathbf{K}_1 through to \mathbf{K}_3 . \mathbf{K}_0 , while not

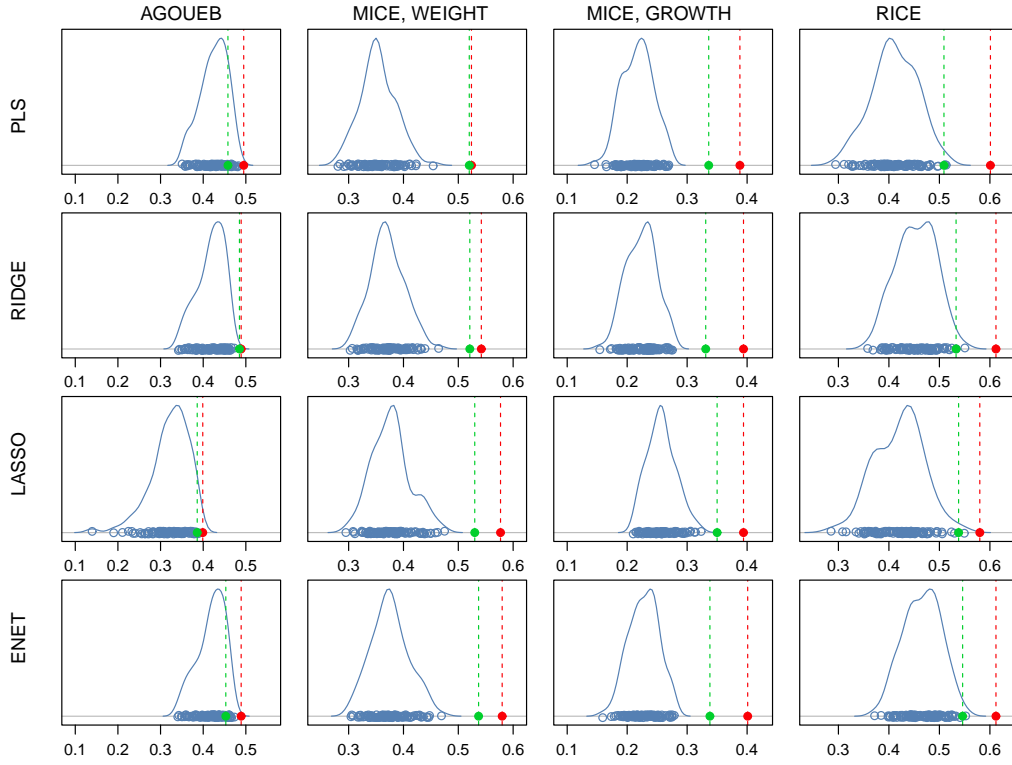


Figure 5: Comparison between the cross-validated correlations obtained from the Markov blankets ($\rho_{CV,MB}$, vertical red dashed line in each panel) and the subsets of the same size obtained from a single-SNP analysis (green dashed line) and from random sampling (blue empirical density curve).

Kinship matrix	$\rho_{CV,MB}$	Δ_1	$\rho_{MB,KIN}$	Δ_2	$\rho_{CV,MB}$	Δ_1	$\rho_{MB,KIN}$	Δ_2
	AGOUEB, YIELD				RICE, SEEDS/PANICLE			
\mathbf{K}_0	0.418	-0.093	0.479	-0.032	0.426	-0.170	0.597	+0.001
\mathbf{K}_1	0.412	-0.100	0.482	-0.030	0.428	-0.161	0.592	+0.002
\mathbf{K}_2	0.414	-0.099	0.491	-0.022	0.429	-0.168	0.589	-0.008
\mathbf{K}_3	0.415	-0.105	0.475	-0.045	0.425	-0.169	0.592	-0.003
	MICE, GROWTH RATE				MICE, WEIGHT			
\mathbf{K}_0	0.194	-0.182	0.378	+0.002	0.219	-0.317	0.534	-0.002
\mathbf{K}_1	0.118	-0.248	0.357	-0.008	0.120	-0.387	0.457	-0.005
\mathbf{K}_2	0.176	-0.211	0.363	-0.025	0.182	-0.326	0.480	-0.028
\mathbf{K}_3	0.195	-0.204	0.379	-0.021	0.225	-0.328	0.530	-0.024

Table 3: Correlation coefficients for GBLUP using the four kinship matrices defined in Section 2.2 and Markov blanket feature selection. $\rho_{CV,MB}$ is defined as in Table 1; $\rho_{MB,KIN}$ is the predictive correlation obtained from cross-validation with the use of Markov blankets but with the kinship matrices estimated from the full SNP profile. The highest value for each quantity and data set is shown in bold. $\Delta_1 = \rho_{CV} - \rho_{CV,MB}$ and $\Delta_2 = \rho_{CV} - \rho_{MB,KIN}$, using the ρ_{CV} reported in Table 2.

being competitive with \mathbf{K}_3 , outperforms at least one of \mathbf{K}_1 and \mathbf{K}_2 for all data sets but AGOUEB. The means of the four ρ_{CV} values are 0.504 for \mathbf{K}_0 , 0.493 for \mathbf{K}_1 , 0.501 for \mathbf{K}_2 , and 0.518 for \mathbf{K}_3 , all with an approximate standard deviation of 0.0057. Thus, GBLUP with \mathbf{K}_3 performs as well as the elastic net and outperforms PLS, ridge regression and the LASSO.

Although the elastic net performed equally well with or without Markov blanket feature selection, that is not the case for GBLUP (Table 3). For all kinship matrices, the reduced size of the Markov blanket relative to the full SNP set detracts from the computation of kinship coefficients, leading to a substantial loss of predictive power. If all SNPs are available and can be used to compute the kinship matrices, then much but not all of this loss is restored.

5 Conclusions

We have shown that Markov blanket feature selection applied as a preliminary step in GS with a continuous trait is able to greatly reduce the size of the SNP set with no loss (and possibly a small gain) in the predictive power of PLS, ridge regression, LASSO and the elastic net. Among those models, the elastic net was the best performer, followed by ridge regression. If GS is to be performed repeatedly

for the same phenotype, for example in successive generations of crops, Markov blanket feature selection opens the possibility of reducing costs by genotyping many fewer markers.

In the absence of a feature selection step, the LD-adjusted kinship matrix \mathbf{K}_3 (Speed et al., 2012a) provides slightly better predictive power than the matrix with no LD adjustment \mathbf{K}_2 (Astle and Balding, 2009) and the IBS kinship matrix \mathbf{K}_0 produced by PLINK (Purcell et al., 2007). In turn, \mathbf{K}_2 and \mathbf{K}_0 appear superior to the matrix with neither LD adjustment nor standardising of SNPs \mathbf{K}_1 (Habier et al., 2007). Using \mathbf{K}_3 , GBLUP was competitive with the elastic net (both had mean $\rho_{CV} = 0.52$ over the four datasets).

Markov blanket feature selection is not compatible with GBLUP because of the requirement for large numbers of SNPs to compute the kinship matrix. However, Markov blanket feature selection has only a small adverse effect on GBLUP if all SNPs are available for computing the kinship matrix.

Acknowledgements

The work presented in this paper forms part of the MIDRIB project, which is funded by the UK Technology Strategy Board (TSB) and Biotechnology & Biological Sciences Research Council (BBSRC), grant TS/I002170/1. We thank our project partners for helpful discussions. We also thank the AGOUEB Consortium (supported by UK DEFRA, the Scottish Government, through the Sustainable Arable LINK Program Grant 302/BB/D522003/1) for making their data available.

References

- Aliferis, C. F., A. Statnikov, I. Tsamardinos, S. Mani, and X. D. Xenofon (2010): “Local Causal and Markov Blanket Induction for Causal Discovery and Feature Selection for Classification Part I: Algorithms and Empirical Evaluation,” *J. Mach. Learn. Res.*, 11, 171–234.
- Astle, W. and D. J. Balding (2009): “Population Structure and Cryptic Relatedness in Genetic Association Studies,” *Stat. Sci.*, 24, 451–471.
- Bravo, H. C., K. E. Leeb, B. E. K. Kleinb, R. Kleinb, S. K. Iyengarc, and G. Wahbad (2009): “Examining the Relative Influence of Familial, Genetic, and Environmental Covariate Information in Flexible Risk Models,” *PNAS*, 106, 8128–8133.

- Cockram, J., J. White, D. L. Zuluaga, D. Smith, J. Comadran, M. Macaulay, Z. Luo, M. J. Kearsey, P. Werner, D. Harrap, C. Tapsell, H. Liu, P. E. Hedley, N. Stein, D. Schulte, B. Steuernagel, D. F. Marshall, W. T. Thomas, L. Ramsay, I. Mackay, D. J. Balding, The AGOUEB Consortium, R. Waugh, and D. M. O’Sullivan (2010): “Genome-Wide Association Mapping to Candidate Polymorphism Resolution in the Unsequenced Barley Genome,” *PNAS*, 107, 21611–21616.
- de los Campos, G., J. M. Hickey, R. Pong-Wong, H. D. Daetwyler, and M. P. L. Calus (2012): “Whole-Genome Regression and Prediction Methods Applied to Plant and Animal Breeding,” *Genetics*, 193, 327–345.
- Forni, S., I. Aguilar, and I. Misztal (2011): “Different Genomic Relationship Matrices for Single-Step Analysis Using Phenotypic, Pedigree and Genomic Information,” *Genet. Sel. Evol.*, 43, 1–7.
- Friedman, J. H., T. Hastie, and R. Tibshirani (2010): “Regularization Paths for Generalized Linear Models via Coordinate Descent,” *J. Stat. Soft.*, 33, 1–22.
- Gianola, D., G. de los Campos, W. G. Hill, E. Manfredi, and R. Fernando (2009): “Additive Genetic Variability and the Bayesian Alphabet,” *Genetics*, 183, 347–363.
- Goeman, J. J. (2012): *penalized R package*, R package version 0.9-41.
- Guan, Y. and M. Stephens (2011): “Bayesian Variable Selection Regression for Genome-Wide Association Studies and Other Large-Scale Problems,” *Ann. Appl. Stat.*, 5, 1780–1815.
- Habier, D., R. L. Fernando, and J. C. M. Dekkers (2007): “The Impact of Genetic Relationship Information on Genome-Assisted Breeding Values,” *Genetics*, 177, 2389–2397.
- Hastie, T., R. Tibshirani, B. Narasimhan, and G. Chu (2012): *impute: Imputation for Microarray Data*, R package version 1.30.0.
- Hayes, B. J., P. J. Bowman, A. J. Chamberlain, and M. E. Goddard (2009): “Genomic Selection in Dairy Cattle: Progress and Challenges,” *J. Dairy Sci.*, 92, 433–443.
- Heffner, E. L., M. E. Sorrells, and J.-L. Jannink (2009): “Genomic Selection for Crop Improvement,” *Crop Sci.*, 49, 1–12.

- Hoerl, A. E. and R. W. Kennard (1970): “Ridge Regression: Biased Estimation for Nonorthogonal Problems,” *Technometrics*, 12, 55–67.
- Hooper, J. W. (1958): “The Sampling Variance of Correlation Coefficients Under Assumptions of Fixed and Mixed Variates,” *Biometrika*, 45, 471–477.
- Hotelling, H. (1953): “New Light on the Correlation Coefficient and Its Transforms,” *J. Roy. Stat. Soc. B*, 15, 193–232.
- Koller, D. and M. Sahami (1996): “Toward optimal feature selection,” in *Proceedings of the 13th International Conference on Machine Learning (ICML)*, 284–292.
- Legendre, P. (2000): “Comparison of Permutation Methods for the Partial Correlation and Partial Mantel Tests,” *J. S. Comput. Sim.*, 67, 37–73.
- Li, Y., C. J. Willer, J. Ding, P. Scheet, and G. R. Abecasis (2010): “MaCH: using Sequence and Genotype Data to Estimate Haplotypes and Unobserved Genotypes,” *Genet. Epidemiol.*, 34, 816–834.
- Macciotta, N. P. P., G. Gaspa, R. Steri, C. Pieramati, P. Carnier, and C. Dimauro (2009): “Pre-Selection of Most Significant SNPs for the estimation of Genomic Breeding Values,” *BMC Proc.*, 3, S14.
- Margaritis, D. (2003): *Learning Bayesian Network Model Structure from Data*, Ph.D. thesis, School of Computer Science, Carnegie-Mellon University, Pittsburgh, PA, available as Technical Report CMU-CS-03-153.
- Meuwissen, T. H. E., B. J. Hayes, and M. E. Goddard (2001): “Prediction of Total Genetic Value Using Genome-Wide Dense Marker Maps,” *Genetics*, 157, 1819–1829.
- Mevik, B.-H., R. Wehrens, and K. H. Liland (2011): *pls: Partial Least Squares and Principal Component Regression*, R package version 2.3-0.
- Morris, A. P. and L. R. Cardon (2007): “Whole Genome Association,” in D. J. Balding, M. Bishop, and C. Cannings, eds., *Handbook of Statistical Genetics*, Wiley, 3rd edition.
- Park, T. and G. Casella (2008): “The Bayesian Lasso,” *J. Am. Stat. Assoc.*, 103, 681–686.
- Pearl, J. (1988): *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann.

- Piepho, H.-P. (2009): “Ridge Regression and Extensions for Genomewide Selection in Maize,” *Crop Sci.*, 49, 1165–1176.
- Piepho, H.-P., J. O. Ogutu, T. Schulz-Streeck, B. Estaghirou, A. Gordillo, and F. Technow (2012): “Efficient Computation of Ridge-Regression Best Linear Unbiased Prediction in Genomic Selection in Plant Breeding,” *Crop Sci.*, 52, 1093–1104.
- Purcell, S., B. Neale, K. Todd-Brown, L. Thomas, M. A. Ferreira, D. Bender, J. Maller, P. Sklar, P. I. de Bakker, M. J. Daly, and P. C. Sham (2007): “PLINK: a Tool Set for Whole-Genome Association and Population-Based Linkage Analyses,” *Am. J. Hum. Genet.*, 81, 559–575.
- Rostoks, N., L. Ramsay, K. MacKenzie, L. Cardle, P. R. Bhat, M. L. Roose, J. T. Svensson, N. Stein, R. K. Varshney, D. F. Marshall, A. Graner, T. J. Close, and R. Waugh (2006): “Recent History of Artificial Outcrossing Facilitates Whole-Genome Association Mapping in Elite Inbred Crop Varieties,” *PNAS*, 106, 18656–18661.
- Schulz-Streeck, T., J. Ogutu, and H.-P. Piepho (2011): “Pre-Selection of Markers for Genomic Selection,” *BMC Proc.*, 5, S12.
- Scutari, M. (2010): “Learning Bayesian Networks with the bnlearn R Package,” *J. Stat. Soft.*, 35, 1–22.
- Scutari, M. and A. Brogini (2012): “Bayesian Network Structure Learning with Permutation Tests,” *Commun. Stat. Theory*, 41, 3233–3243, special Issue “Statistics for Complex Problems: Permutation Testing Methods and Related Topics”. Proceedings of the Conference “Statistics for Complex Problems: the Multivariate Permutation Approach and Related Topics”, Padova, June 14–15, 2010.
- Solberg, L. C., W. Valdar, D. Gauguier, G. Nunez, A. Taylor, S. Burnett, C. Arboledas-Hita, P. Hernandez-Pliego, S. Davidson, P. Burns, S. Bhattacharya, T. Hough, D. Higgs, P. K. W. O. Cookson, Y. Zhang, R. M. Deacon, J. N. Rawlins, R. Mott, and J. Flint (2006): “A protocol for high-throughput phenotyping, suitable for quantitative trait analysis in mice,” *Mamm. Genome*, 17, 129–146.
- Speed, D., G. Hermani, M. R. Johnson, and D. J. Balding (2012a): “Improved Heritability Estimation from Genome-Wide SNPs,” *Am. J. Hum. Genet.*, 91, 1011–1021.

- Speed, D., G. Hermani, M. R. Johnson, and D. J. Balding (2012b): *LDAK*, <http://dougspeed.com/ldak/>.
- Tibshirani, R. (1996): “Regression Shrinkage and Selection via the Lasso,” *J. Roy. Stat. Soc. B*, 58, 267–288.
- Tsamardinos, I., C. F. Aliferis, and A. Statnikov (2003): “Algorithms for Large Scale Markov Blanket Discovery,” in *Proceedings of the 16th International Florida Artificial Intelligence Research Society Conference*, 376–381.
- Valdar, W., L. C. Solberg, D. Gauguier, S. Burnett, P. Klenerman, W. O. Cookson, M. S. Taylor, J. N. Rawlins, R. Mott, and J. Flint (2006): “Genome-Wide Genetic Association of Complex Traits in Heterogeneous Stock Mice,” *Nat. Genet.*, 8, 879–887.
- VanRaden, P. (2008): “Efficient Methods to Compute Genomic Predictions,” *J. Dairy Sci.*, 91, 4414–4423.
- Vazquez, A. I., G. de los Campos, Y. C. Klimentidis, G. J. M. Rosa, D. Gianola, N. Yi, and D. B. Allison (2012): “A comprehensive genetic approach for improving prediction of skin cancer risk in humans,” *Genetics*, 192, 1493–1502.
- Waugh, R., D. Marshall, B. Thomas, J. Comadran, J. Russell, T. Close, N. Stein, P. Hayes, G. Muehlbauer, J. Cockram, D. O’Sullivan, I. Mackay, A. Flavell, AGOUEB, BarleyCAP, and L. Ramsay (2010): “Whole-Genome Association Mapping in Elite Inbred Crop Varieties,” *Genome*, 53, 967–972.
- Wimmer, V., T. Albrecht, H.-J. Auinger, and C.-C. Schön (2012): “synbreed: Framework for the Analysis of Genomic Prediction Data Using R,” *Bioinformatics*, 18, 2086–2087.
- Zhao, K., C. Tung, G. C. Eizenga, M. H. Wright, M. L. Ali, A. H. Price, G. J. Norton, M. R. Islam, A. Reynolds, J. Mezey, A. M. McClung, C. D. Bustamante, and S. R. McCouch (2011): “Genome-Wide Association Mapping Reveals a Rich Genetic Architecture of Complex Traits in *Oryza Sativa*,” *Nat. Commun.*, 2, 467.
- Zou, H. and T. Hastie (2005): “Regularization and Variable Selection via the Elastic Net,” *J. Roy. Stat. Soc. B*, 67, 301–320.